DOCUMENT RESUME

ED 413 361 TM 027 700

AUTHOR McGinty, Dixie; Neel, John H.
TITLE Judgmental Standard Setting Using a Cognitive Components Model.
PUB DATE 1996-04-00
NOTE 23p.; Paper presented at the Annual Meeting of the National Council on Measurement in Education (New York, NY, April 4-8, 1996). For a replication study, see TM 027 701.
PUB TYPE Reports - Evaluative (142) -- Speeches/Meeting Papers (150)
EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS Cognitive Processes; Criterion Referenced Tests; Evaluation Methods; Grade 3; *Interrater Reliability; *Judges; *Mathematics Tests; Primary Education; Scoring; Standards; *Test Items
IDENTIFIERS Angoff Methods; *Cognitive Component Analysis; *Standard Setting

ABSTRACT

A new standard setting approach is introduced, called the cognitive components approach. Like the Angoff method, the cognitive components method generates minimum pass levels (MPLs) for each item. In both approaches, the item MPLs are summed for each judge, then averaged across judges to yield the standard. In the cognitive components approach, items must be decomposed into nonoverlapping cognitive components that may be thought of as specific skills or knowledge required for a correct response to an item. The method was studied with 12 judges, all third- or fourth-grade teachers who judged a sample of 2,500 students from a third-grade state criterion-referenced mathematics test. Teachers also used the Angoff method to set standards for these results. The most surprising finding of the study was the similarity between the two sets of results. Results from the cognitive component method resembled those from the Angoff method in the range and standard deviation of the recommended standards, as well as in the final standard itself. Interjudge variability was considerably smaller for the cognitive components responses than for the Angoff responses. Some of the validity concerns that may be raised by the cognitive components method are discussed. Additional studies are necessary to support the use of the method in setting standards, and the method is probably only useful when test items lend themselves to decomposition into subtasks. (Contains 3 tables and 10 references.) (SLD)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED 413 361

Judgmental Standard Setting Using a Cognitive Components Model

Dixie McGinty

John H. Neel

Georgia State University

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☑ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL HAS BEEN GRANTED BY

Dixie McGinty

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

TM027700

Paper presented at the 1996 annual meeting of the National Council on Measurement in Education, New York

BEST COPY AVAILABLE

Introduction

Despite several decades of research, setting standards for minimal competency tests remains problematic. The popular judgmental methods proposed by Angoff (1971), Ebel (1972), Jaeger (1989) and Nedelsky (1954) have much to recommend them, and admirable efforts continue to be made to refine these processes. Nevertheless, standard-setting procedures continue to be fraught with difficulties that make them vulnerable to harsh criticism. One of the most salient problems is that the recommendations of the judges are often substantially more variable than might be hoped, which reduces our confidence in the standard that has been set.

Variability in judges' recommendations for a set of items can result from either (1) differing opinions about what should be required of examinees or (2) differing perceptions of the test items and the cognitive demands they pose. We suggest that the first type of variability is to be expected; we expect individuals to differ in their opinions and in the value judgments they make. The second type of variability, on the other hand, is potentially more threatening; it results from judges' varying abilities to perceive correctly the important features of test items.

Our research was motivated by the belief that the judges' task in standard setting can, and should be, made easier, especially with regard to the perception of items. Judges, even those who are teachers, typically have limited experience with

actual test items, and many lack training in cognitive psychology. It seems unrealistic to expect them to become skillful in assessing the difficulty of items after just a few hours of training. Further, the time allowed for standard setting may not be sufficient for judges to thoroughly analyze each item and identify the skills it demands. We thus believe that a new method is needed, one that takes some of the guesswork out of the prediction of item difficulty.

In this study, we introduce a new standard-setting approach, which we call the cognitive components approach. Like the Angoff method, the cognitive components method generates minimum pass levels (MPLs) for each item. In both approaches, the item MPLs are summed for each judge, then averaged across judges to yield the standard. What makes the cognitive components approach different is that the MPLs are arrived at in a different, less direct manner. Instead of making judgments about the test items directly, judges specify minimum success rates for item subtasks or components; these values are then put together to form a synthetic MPL for each item.

Before judges convene, items must be decomposed into nonoverlapping "cognitive components," which may be thought of as specific skills, subtasks, or pieces of knowledge that are believed to be required for a correct response to an item. Consider, for example, the following estimation item, which is similar to one item on the test used in this study (assume that the response options are all multiples of 100):

"516 + 193 + 232 is about _____."

One way to decompose this item is to posit that, in order to respond correctly, an examinee must (1) recognize "about" as a prompt for rounding or estimation, (2) round three-digit numbers to the nearest hundred, (3) recognize "+" as a prompt for addition, (4) line numbers up vertically for addition, and (5) apply basic addition facts.

When the judges convene, they are not presented with actual test items, but rather with brief statements or descriptions of cognitive components. For each, the judges are asked to complete a statement of the type, "In order to pass the test, an examinee must be able to apply this skill correctly at least _____% of the time." In other words, judges specify the minimum ratio of the number of correct applications of the specific component to the number of situations that require it (note that this is not equivalent to the percent of <u>items</u> requiring that skill which should be answered correctly). We call this value the minimum success rate (MSR) for the cognitive component; it is equivalent to the probability that a minimally qualified examinee will apply the cognitive component successfully.

To compute the minimum pass level for a given item, the minimum success rates (MSRs) for those cognitive components that have been identified for that item are simply multiplied together. For example, in this study, the five components listed above for the item shown had average MSRs of .775, .667, .996, .883, and .973, respectively; thus the synthetic MPL for this

hypothetical item would be (.775)(.667)(.996)(.883)(.973), or .4423. Clearly, an assumption of independence of components underlies this model. As Embretson (1984) notes, empirical support for the notion of independent components that contribute additively to task difficulty can be found in the work of Sternberg (1969) and Pachella (1974). We do feel, however, that the consequences of this assumption must be carefully considered in order to assure the validity of the standard-setting process. The independence assumption will thus be discussed further in a later section of this paper.

Methods

In this study, we tried out the cognitive components model in an experimental setting and compared its results to those yielded by the Angoff model. Our purpose was an initial, exploratory investigation of the new approach to determine whether it is worthy of further study. We were interested, for example, in how judges would perceive the task of making judgments about cognitive components rather than about items. We also wondered about the range and variability of the probability values they would assign to the components, and about how the resulting minimum pass levels (MPLs) for items would compare to those generated by the Angoff method, as well as to empirical item difficulties.

Twelve judges, all of whom were third- or fourth-grade teachers at the time of their participation in the study, completed a simulated standard-setting exercise in which both the

cognitive components model and the Angoff model were used to arrive at two separate standards for the same set of items. The items used were a subset of 55 items from the mathematics portion of a statewide criterion-referenced test for third grade, which was used to make decisions about promotion of students to fourth grade. Items were multiple choice with four options each. A variety of item types were represented, e.g. computation, estimation, simple word problems, computation with money, reading tables, etc.

Insert Table 1 about here

The test was administered in 1991 to a population of approximately 90,000 third graders statewide. Empirical data on the items was available based on a systematic sample of 2500 students. P-values for the items used in this study ranged from .62 to .99, with a mean of .85.

Before the judges met, items were decomposed into cognitive components. The resulting set of 29 components is given in Table 1. While we considered this set to be workable for an initial study, it should by no means be viewed as definitive, but rather as preliminary. The number of components represented by each item ranged from 1 to 8, with only a few items representing only one. Items for which only one component was identified were those whose type and format differed markedly from the remaining items, making the already-identified components inapplicable.

While these items probably could have also been decomposed in some way, it did not seem to be worth our time to do so for the purposes of this initial study. We must emphasize that the set of components we used suffices only for a preliminary look at the feasibility of this type of standard-setting approach.

Judges completed the exercise in groups of three on each of four different occasions. Two groups of three provided ratings using the Angoff method first, followed by the cognitive components method; the other two groups applied the two methods in the reverse order. Assignment to either the Angoff-first or the cognitive-components first condition was done randomly. All groups of judges received extensive instructions for both methods based on a prepared script. In addition to specifying probabilities, for each item and cognitive component judges were also asked to respond to a confidence item ("How confident do you feel about your response?") using a 5-point Likert scale. After providing both sets of ratings, judges were asked several open-ended questions about how they perceived the two methods, and these were discussed by each group.

Results

The cognitive components method resulted in a final standard of 65.6%, or 36.1 items correct, while the Angoff method yielded a somewhat lower standard of 58.8%, or 32.3 items correct. Standards recommended by individual judges are given in Table 2. A surprising result is that the minimum and maximum recommended standards were virtually identical across methods, despite the

BEST COPY AVAILABLE

fact that the judges used two very different thought processes. Examination of the raw response data for Judge 2 reveals that this judge tended to assign extremely high values, very often 1.00, across the board to both items and cognitive components, which makes the congruence of the maximum standard across methods less surprising. No pattern, however, is immediately discernible in the raw responses of Judge 8, who provided the minimum standard in both instances. While the correlation between the two sets of recommended standards was .63 ($.01 < p < .05$), the relationship between the sets of item MPLs generated by the two methods was more impressive ($r = 0.59$, $p < .0001$). It should be noted that any agreement between the results yielded by the two methods should be interpreted somewhat cautiously since judges' exposure to the first method they used may have influenced their responses using the second method.

Insert Table 2 about here

With regard to interjudge variability, our results are encouraging in some ways and discouraging in others. The use of the cognitive components method did not result in lower variability of the final standard; standard deviations of the standard were 9.24 for the cognitive components method and 9.07 for the Angoff method. Individual item MPLs, however, were less variable across judges with the cognitive components method than with the Angoff method. Standard deviations of the MPLs

generated by the cognitive components method were lower than those resulting from the Angoff method for 38 out of 55 items ($p < .01$); mean MPL standard deviations were .1924 and .2366 for the cognitive components method and the Angoff method, respectively. Consideration of the raw response data (i.e., component MSRs for the cognitive components model, item MPLs for the Angoff method) suggests that these judges tended to agree substantially more about the cognitive components that make up the items than about the items themselves. The mean standard deviation for cognitive components across judges was .1453, as compared to .2366 for items using the Angoff method. Intercorrelations among judges are also revealing. Out of the 66 possible correlations, the MSR-level responses generated by the cognitive components model resulted in 47 correlations that were significant at the .01 level, while the Angoff data resulted in only 29 that were significant at that level.

The two methods were similar in the degree to which item MPLs were correlated with empirical item p-values. Correlations were .63 for the Angoff method ($p < .0001$) and .57 for the cognitive components method ($p < .0001$). Though the Angoff method fared somewhat better in this regard, we are encouraged by the result for the new method, especially since the specific set of cognitive components used was extremely preliminary. A more refined set of components would hopefully lead to an even greater correlation between MPLs and item difficulty. We must add, however, that we do not feel that extremely high correlations are

BEST COPY AVAILABLE

necessary to ensure validity of a standard-setting process. Positive correlations support validity because they suggest that, in making their recommendations, judges are attending to those features of the items that actually contribute to item difficulty. At the same time, the judges' task is not equivalent to a simple prediction of item difficulty, but instead requires both accurate perception of item features and value judgments about the importance of specific items or skills; thus a very strong correlation might also be suspect.

Responses to the confidence-level question for each item and cognitive component suggest that judges felt more confident about the judgments they made about cognitive components than about the Angoff probabilities they specified for each item. Mean confidence ratings across items (for the Angoff method) or cognitive components (for the cognitive components method) are given in Table 3. For every judge, the mean rating was higher for the cognitive components method than for the Angoff method; for eight of the twelve judges, these differences were

---

Insert Table 3 about here

---

significant at the .01 level. An open-ended discussion following the standard-setting exercises underscored these results: Almost unanimously, the judges said they preferred the cognitive components method and found it easier. Many of them reported difficulty in conceptualizing a "minimally competent" examinee,

and several suggested that dealing with an item was more complicated than dealing with a skill or cognitive component. When asked specifically about the relative confidence they would have in an actual standard set by the each of the two methods, most still preferred the cognitive components method, though two judges changed their preference to the Angoff method, noting that the formats of the items should be considered in setting a standard. Even after this point had been raised, however, the remaining judges stood by the cognitive components method. As one teacher said, "This method fits the way we teachers think." Despite the acclaim found in this study for the new method, however, we question whether judges would really feel the same way if this were an actual standard-setting situation.

Discussion

The most surprising finding of our study was the similarity between the two sets of results. Results from the cognitive components method resembled those from the Angoff method in the range and standard deviation of the recommended standards, as well as in the final standard itself. Though we reject the notion that the Angoff method should be the measuring stick by which other methods are evaluated, we must admit that we find this result somewhat reassuring; we are relieved that, at least in this one instance, our new method did not lead to preposterous results. Clearly, however, the study needs to be replicated a number of times before conclusions can be drawn, and some of these replications should employ two different, though randomly

equivalent, sets of judges in order to ensure that the two methods do not contaminate each other.

Another result that we find important is that interjudge variability was considerably smaller for the cognitive components raw responses (i.e., for the specification of minimum success rates for each component) than for the Angoff responses. In other words, our judges tended to agree more strongly about the components that compose the items than they did about the items themselves. This seems to suggest to us that there was considerable agreement among judges as to what should be required, but that judges differed in how they perceived the actual items. It could be argued, we think, that the relative contributions of value judgment and perception of difficulty to the judges' task differ across the two methods. More specifically, it seems to us that in specifying a minimum success rate for a cognitive component, value judgment plays a relatively greater role, and perception a relatively smaller role, than in specifying an item MPL using the Angoff approach. If future studies of the cognitive components model yield similar results, this may suggest that teachers' value judgments do not vary as greatly as one might think, which is an encouraging prospect for standard setting in general.

The cognitive components model, if used as a self-standing standard-setting procedure, raises some validity concerns that are very different from those posed by the Angoff procedure. Any set of items can, of course, be decomposed into cognitive

BEST COPY AVAILABLE

components in a wide variety of ways. With regard to the estimation item presented earlier in this paper, one could argue that the component "round three-digit numbers..." encompasses several steps, each of which could have been named as a separate cognitive component. Further, it is also possible to argue that " recognize '+' " should be subsumed under "apply basic addition facts." The outcome of the standard-setting process is likely to differ when applied to different sets of components, and, since there is no one "correct" set of components, special efforts must be taken to bolster the validity of the process.

First, we recommend that the cognitive components be subject to a validation procedure by groups of educators and psychologists, similar to the processes often used at the time of test development to validate the content of high-stakes tests. The central question to be addressed by such a procedure is whether the set of components fairly represents the items on the test. While this validation process could be performed post hoc on an existing test before setting standards, it fits in very nicely at the test development stage, since items can then be developed and tests constructed to match a precise configuration of the desired components.

Second, if field-test data are available for the items, we suggest using multiple regression analysis to determine how well the combinations of components identified for each item account for the variability in empirical item difficulties. In other words, each cognitive component is represented as a dichotomous

categorical variable whose value is 1 for items that involve that component and 0 for items that do not; item p-values are then regressed on the resulting vectors of binary data. This approach has been used by several researchers, e.g. Tatsuoka (1990). While a high coefficient of determination (r^2) does not guarantee that the set of components is "correct," it does suggest that the item features represented by the components are among those that contribute to item difficulty. We recommend an iterative process for identifying and validating sets of components, i.e., one which involves identifying a preliminary set, performing the regression analysis, and revising the set until a higher r^2 is obtained. This goal must be balanced, however, against the need to end up with a manageable number of components, all of which can be communicated clearly to judges.

If the cognitive components model is to be pursued further, the assumption that the cognitive components are independent must be examined closely, both empirically and logically, and the consequences of violating it must be considered. A faulty assumption may compromise the validity of the interpretation given to the resulting standard. On the other hand, many measurement procedures are based on assumptions that are not likely to be met in reality. Our opinion, then, is that it is too early to dismiss the cognitive components model on the basis of this issue alone.

The real contribution of the cognitive components model may lie in its potential for combination with other judgmental

standard-setting methods. Several intriguing possibilities seem to warrant research. For example, a measure of intrajudge consistency could be devised that would involve comparing each judge's Angoff ratings to the MPLs generated for the same items using the cognitive components method. This measure would offer some advantages over the IRT-based measure of intrajudge consistency proposed by Van der Linden (1982). First, it would in no way depend on empirical item data, which is an appealing prospect since the judges' task in standard setting should not be reduced to simply predicting item difficulty, though that it is part of it. Second, the use of IRT-based indices with standard-setting data is somewhat problematic since the assumptions of IRT do not necessarily hold for standard-setting data.

In another possible application of the cognitive components model, judges would rate cognitive components as a preliminary step before seeing the actual test items, then provide Angoff item ratings in the normal manner. While the actual standard would be computed using the Angoff model, the cognitive components data collected earlier could be presented to the judges for their consideration in revising their original Angoff responses. In other words, each judge would be shown the synthetic MPLs that resulted from his/her own responses to the cognitive components, along with the component-level data that generated them. This provides an interesting alternative to the use of normative data as a supplement to the Angoff process;

however, there is no reason why normative data could not still be used along with the cognitive-components data.

If a standard is to be computed using the cognitive components model as the primary method, we recommend that judges be allowed to see the actual test items at some point, possibly in a process resembling a reversal of the one described above. The more information available to aid judges, the more sound we can expect the resulting decision to be. If items are available, and if time permits, we see no reason to "hide" them from the judges.

Conclusion

Much more research is needed in order to conclude that the cognitive components model offers a viable approach. Even in the presence of more empirical evidence in its support, the model raises many issues that would need to be discussed and argued a priori. We do feel, however, that the results of this initial study are certainly not discouraging. Clearly, our current model is just a starting point for a new approach, and it needs refinements. For example, we are currently considering adjusting the model to account for guessing, and other challenges lie ahead as well.

Standard-setting procedures will always be imperfect due to the nature of the task they are intended to accomplish. While some methods may be clearly superior to others, the choice among methods may often be a question of which specific methodological weaknesses one is willing to accept and of what trade-offs one is

willing to make. Further, some methods may be more appropriate for certain types of tests or judges than for others.

The cognitive components model would appear at this point to be applicable only to those types of tests whose items lend themselves to decomposition into subtasks. On the other hand, due to rapid advances in cognitive psychology and artificial intelligence, we may soon find that this approach or a similar one is applicable to more types of tests than we had originally thought. Perhaps similar but more sophisticated models could be developed that incorporate not only what we call cognitive components, but other item characteristics as well.

We would welcome comments or suggestions, preferably via e-mail, from anyone who would like to make them. Dixie's address is <epsdlmx@gsusgi2.gsu.edu>; John's is <jneel@gsu.edu>. You can also write to either of us c/o Department of Educational Policy Studies, University Plaza, Atlanta, GA 30303-3083.

Table 1

Cognitive Components Used in this Study

C2. Translate words to numerals.
C3. Choose the correct operation to solve a word problem.
C5. Count objects in a picture.
C6. Understand what is meant by "tens" and "ones" in place value.
C7. Compare two numbers to determine which is greater.
C8. Apply basic addition facts.
C9. Line up amounts of money vertically for computation.
C10. Regroup (in addition).
C11. Recognize "+" as a prompt for addition.
C12. Compare three or more numbers to determine which is greatest.
C13. Read a table.
C14. Know the monetary value of a pictured coin.
C15. In a subtraction word problem, know which number to subtract from which.
C16. Apply basic subtraction facts.
C17. Select an appropriate unit of measure.
C18. Recognize "-" as a prompt for subtraction.
C19. Round three-digit numbers to the nearest hundred.
C20. Line up two- or three-digit numbers vertically for computation.
C24. Compare sizes of pictured objects.
C25. Recognize "about" as a prompt for estimation or rounding.
C27. Know what is meant by perimeter of a figure.
C29. Regroup (in subtraction).
C33. Read a bar graph.
C38. Recognize "x" as a prompt for multiplication.
C39. Apply basic multiplication facts.
C46. Recognize ÷ as a prompt for division.
C47. Apply basic division facts.
C49. Recognize ⟌ as a prompt for division.
C53. Know the monetary value of a coin by its name.

Table 2
Recommended Standards by Judge
(Number of Items Correct)

	Angoff Method	Cognitive Components Method
Judge 1	24.10	35.71
Judge 2	50.05	49.98
Judge 3	30.25	44.14
Judge 4	30.05	20.59
Judge 5	37.65	38.11
Judge 6	29.31	27.33
Judge 7	25.28	40.26
Judge 8	20.00	20.43
Judge 9	23.40	34.16
Judge 10	37.11	38.12
Judge 11	35.69	43.96
Judge 12	45.03	40.90
MEAN	32.33	36.14
ST DEV	9.07	9.24

Table 3
Confidence Level Data
Mean Ratings

	Angoff Method (n = 54*)		Cognitive Components Method (n = 29)		F	p > F
Judge	Mean	St Dev	Mean	St Dev		
1	3.643	0.673	4.138	0.743	9.64	.0026
2	3.036	0.267	4.345	0.484	259.31	.0001
3	3.768	0.504	4.621	0.494	55.43	.0001
4	3.196	0.796	3.862	0.875	12.48	.0007
5	3.821	0.834	3.897	0.772	0.16	.6874
6	3.571	0.499	3.897	0.310	10.22	.0020
7	3.218	0.417	4.379	0.494	129.51	.0001
8	3.167	0.458	3.379	0.728	2.87	.0939
9	3.418	0.498	3.759	0.912	4.92	.0293
10	3.696	0.464	3.897	0.724	2.39	.1256
11	3.661	0.668	4.379	0.677	21.91	.0001
12	4.071	0.828	4.793	0.491	18.57	.0001

*One item was inadvertently omitted from this analysis.

References

Angoff, W. H. (1971). Scales,norms, and equivalent scores. In R. L. Thorndike (Ed.), Educational measurement (2nd ed., pp. 508-600). Washington, DC: American Council on Education.

Ebel, R. L. (1972). Essentials of educational measurement. Englewood Cliffs, NJ: Prentice Hall.

Educational Testing Service [ETS] (1976). Report on the study of the use of the National Teachers Examination by the State of South Carolina. Princeton, NJ: ETS.

Embretson, S. E. (1984). A general latent trait model for response processes. Psychometrika, 49, 175-186.

Jaeger, R. M. (1989). Certification of student competence. In R. L. Linn (Ed.), Educational measurement (3rd. ed., pp. 485-514). New York: American Council on Education and Macmillan.

Nedelsky, L. (1954). Absolute grading standards for objective tests. Educational and Psychological Measurement, 14, 3-19.

Pachella, R. G. (1974). The interpretation of reaction time in information-processing research. In B. Kantowitz (Ed.), Human information processing: Tutorials in performance and cognition (pp.41-82). Hillsdale, N.J.:Erlbaum.

Sternberg, R. S. (1969). Memory-scanning: Mental processes revealed by reaction time experiments. American Scientist, 4, 412-457.

Tatsuoka, K. (1990). Toward an integration of item response theory and cognitive error diagnosis. In Frederiksen, N.,

BEST COPY AVAILABLE

Glaser, R., Lesgold, A., & Shafto, M. G. (Eds.), Diagnostic monitoring of skill and knowledge acquisition. Hillsdale, NJ: Erlbaum.

Van der Linden, W. J. (1982). A latent trait method for determining intrajudge consistency in the Angoff and Nedelsky techniques of standard setting. Journal of Educational Measurement, 19(4), 295-308.